Simple networks for spike-timing based computation

Carlos D. Brody
Cold Spring Harbor Laboratory
PO Box 100, Cold Spring Harbor NY 11724

J. J. Hopfield
Department of Molecular Biology
Princeton University, Princeton NJ 08544-1014

Abstract

Spike timing synchronization across different neurons can be selective for the situation where the neurons are driven at similar firing rates, a "many are equals" computation. Such synchronization can be achieved, in the absence of synaptic interactions between the synchronizing neurons, through phase-locking to a common underlying oscillatory potential. Based on this principle, we instantiate an abstract algorithm for robust odor recognition into a model network of spiking neurons whose main features are taken directly from known properties of the primary stages of biological olfactory systems. In this network, recognition of known odors is signaled by spike timing synchronization of specific subsets of 'mitral cells.' This synchronization is highly odor-selective and invariant to a wide range of odor concentrations. It is also robust to the presence of strong distractor odors, thus allowing odor segmentation within complex olfactory scenes. Information about odors is encoded in both the identity of glomeruli activated above threshold (1 bit of information per glomerulus), and in the analog degree of activation of the glomeruli (approximately 3 bits per glomerulus).

Introduction

We have previously described a spiking neural network model in which a computation dubbed 'MANY ARE EQUALS' was imbedded (Hopfield and Brody, 2000; 2001). In this network, when the external inputs to a large group of neurons were equal or closely similar to each other, synaptic interactions between the neurons were such that the neurons' action potentials became synchronized. Thus, synchronization served to compute, and indicate, similarity of inputs. The computational power of this operation was demonstrated by using it to construct a simple model network that was capable of a complex computation, recognizing a short spoken word. Appropriate invariance is a feature of many pattern recognition problems. The 'MANY ARE EQUALS' computation contains an in-built invariance in that synchronization does not depend on the particular input level that is similar across neurons. In our example network, this invariance was used to achieve time-warp invariant recognition.

We now describe how the 'MANY ARE EQUALS' computation can be implemented in a greatly simplified network in which synaptic interactions between the synchronizing neurons are not necessary. The lateral synaptic interactions between neurons, which earlier were responsible for the synchronization phenomenon, are instead abstracted into a common subthreshold oscillation that provides a common timing signal to all the neurons. Synchronization of action potentials across neurons can then be achieved through each neuron's phase-locking its action potentials to the common oscillation; when the neurons phase lock at the same phase, their input currents must be equal. We further show that the problem of intensity-independent odor recognition can be transformed into a 'MANY ARE EQUALS'problem. Using the abstract sensory model described by Hopfield (1999), we show that this implementation of odor recognition by a network of spiking neurons can deal with mixtures and unknown backgrounds as well as a wide range of concentrations.

Our new network, although simplified, still requires recognition units that will detect synchrony in selected subsets of sensory neurons. (Such recognition units will be neurons selective for words in the audition case, and highly selective for odors in the olfaction case). The identity of the sensory stimulus that recognition units are selective for is encoded in the pattern of feedforward synaptic connections from the sensory units to the recognition unit. A biological system must learn such connections, and in a companion paper (Hopfield and Brody, following paper) we develop the synaptic plasticity rules that can lead to the development of such feedforward connections in a biologically plausible fashion.

Our approach here and in the companion paper has not been to develop highly detailed models that mimic all features of neurobiological systems. Rather, we have tried to develop models that most clearly and simply illustrate the essence of the computations. Many features of the models have been directly inspired by properties of biological olfactory systems. In particular, the basic architecture (e.g., input dimensionality, convergence/divergence) has been chosen to match that of the mammalian olfactory bulb. As a shorthand for communication, we therefore sometimes use biological terms to refer to matching concepts in the model (e.g., "glomeruli" encode the activity of a single receptor type, and each glomerulus provides the principal input to a number of "mitral cells"). When successful, such simplified models can be used as guides for later development of more detailed and biological models. For example, the common subthreshold oscillatory drive, observed in the primary stages of biological olfactory systems, and used here, could be created in a variety of ways: through an external oscillator, through an embedded oscillating network, or through feedback synaptic connections between the synchronizing neurons themselves. However, elucidation of the specific mechanism responsible for such an oscillation is not necessary at this computational level of analysis.

The basic synchronization phenomenon

A non-adapting neuron that is driven by an oscillating subthreshold potential plus a constant current input can phase-lock its action potentials with the oscillating potential. This phenomenon is responsible for the phase-locking of the action potentials coming from the cochlear nucleus neurons in response to low frequency tones (Lavine (1971); Johnson (1980)). It is a phenomenon of most neural models, including Hodgkin-Huxley neurons (simulation not shown) and integrate-and-fire neurons. There is a range of strengths of the constant current input at which the phaselocking will be 1-to-1 in the sense that each cycle of the oscillation will contain one spike from the neuron. In the absence of noise, the phase of this spike with respect to the underlying oscillation will be precisely determined by the strength of the constant current input.

In the presence of noise, the phasing will no longer be precise, but there will nevertheless be a tendency to fire at a specific phase. Fig.(1) shows the phase locking of integrate-and-fire neurons having different current inputs in the presence of a common sinusoidal input current.

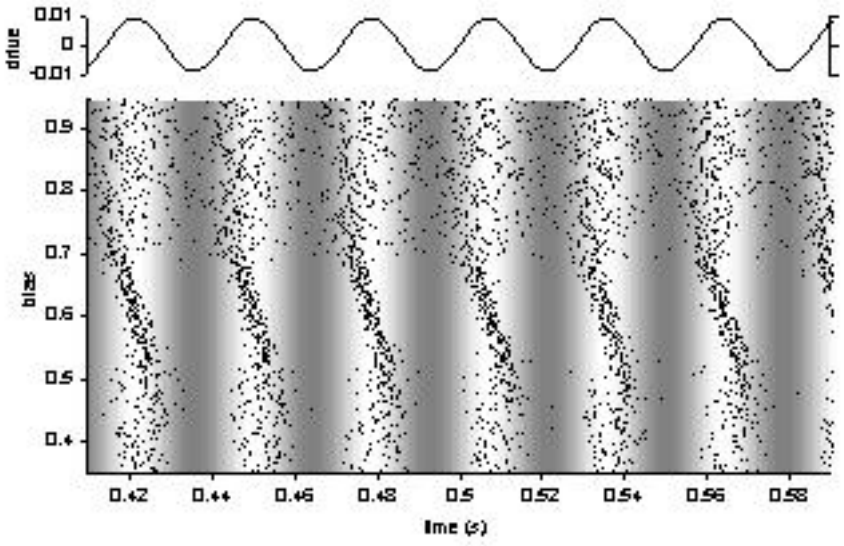

Figure 1. Action potential rasters for integrate-and-fire neurons in the presence of a subthreshold sinusoidal current injection. Each row represents a neuron, and each neuron has a different DC current injected into it, indicated by the ordinate. Dots indicate the time of action potentials. The common sinusoidal input current is indicated in the upper plot, and by the gray-scale shading underlying the action potential rasters. Gaussian noise is injected into each cell. DC currents between 0.5 and 0.7 (arbitrary units) result in excellent phase-locking, with phase progressing in an orderly and almost linear fashion along this current range. The cell time constant was 20 msec, and the drive frequency 35 Hz.

These phase-locking properties can be used to construct an MAE operation. If many neurons are all receiving a current input that is within the 1-to-1 phaselocking range, and all of these input currents are similar, then the neurons will all be firing at a similar phase with respect to the underlying oscillation. The neurons will therefore all be synchronized to each other. In contrast, if the neurons receive constant current inputs that are different from each other, or that are not within the 1-to-1 phaselocking range, then the spike timing synchronization across neurons will

be much weaker. Thus, as in our previous network (Hopfield and Brody, 2000, 2001) similarity of inputs is signaled by synchronization of action potentials. In the present network, however, synchronization across neurons is achieved through synchronization to a common underlying signal, instead of being achieved through direct neuron-to-neuron synaptic interactions.

A comparison of the MAE operation, computed through two different synchronization-promoting mechanisms, is illustrated in Fig.(2). The left hand half of this figure illustrates phase locking via 'horizontal' connections as presented in earlier work; the right hand the same MAE operation computed by a network without horizontal connections, but using instead a common input sub-threshold current to produce synchronization. Similar synchronization is seen with both mechanisms, although there are differences. For example, the interspike separation when well synchronized is always the reciprocal of the period of the common drive in one case (right column), while it depends on the current level at which the currents converge in the other (left column).

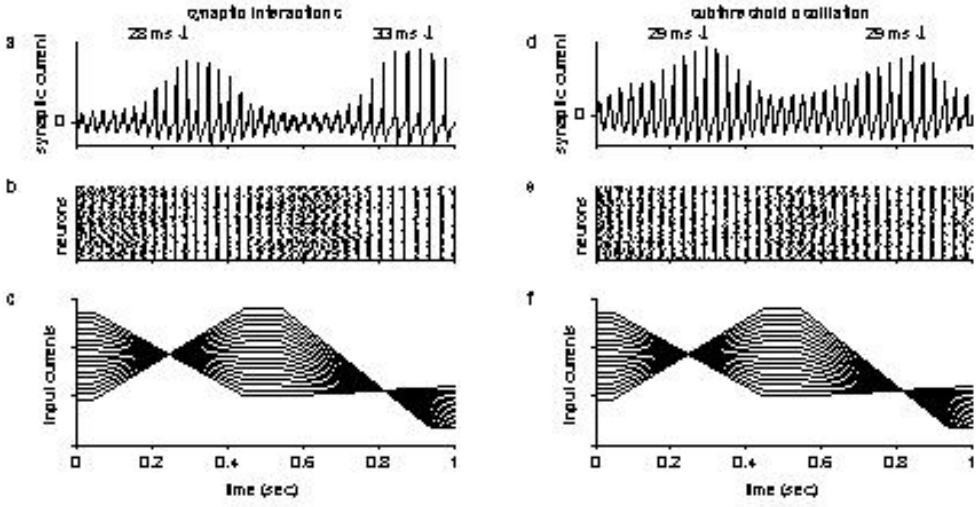

Figure 2. "Many are equals" synchrony, indicating that the input currents to many cells are about the same, implemented through two different mechanisms: 'horizontal' synaptic interactions between a set of cells (a-c), versus a common input sinusoid to otherwise independent cells (d-f). Each cell in panels (b) and (e) has a time-dependent current injected into it; the currents for these cells are shown in (c) and (f). The spike rasters and their movement into and out of synchronization are shown in (b) and (e). The synchrony is also illustrated in panels (a) and (d), which show the membrane potential of a neuron that receives both excitatory (synaptic currents exponential, τ = 2 ms) and inhibitory (alpha function currents, τ = 6 ms.) input currents from the spike rasters shown in (b) and (e), respectively. For panels (a) and (d), the cell time constant is 6 ms. In (d-e), the underlying input frequency is 35 Hz.

Odor recognition

Consider a rat inhabitant of the New York City subway, scurrying along the tracks in search of food. For olfaction to be a useful modality in this endeavor, the rat must be able to *recognize* **odors. But food odors in the natural world will rarely exist alone—they will usually be experienced in the context of a large variety of other odors also present in the environment. Thus, the rat must be able to** *segment***, or separate, a known odor from its olfactory background (similar to the way an experienced cook may identify, by smell, a particular spice used in a soup). Furthermore, the rat will experience an odor at various distances from the source, implying experience of it at many different concentrations. Yet all these different concentrations must be interpreted as coming from the same odor source; odor recognition must thus be significantly** *concentration-invariant***.**

Hopfield (1999) proposed an abstract algorithm that simultaneously addressed these three problems. The essential idea is based on the fact that there is a large family of odor receptors cell types, numbering approximately 1,000. Each receptor cell class responds to many different odors (Sicard and Holley 1984; Buck 1996) and any particular odor thus activates, in a concentration-dependent manner, a substantial subset of these receptors (on the order of hundreds of receptor cell types). In the rat, each glomerulus in the olfactory bulb receives input from ~ 10,000 sensory cells (Shepherd and Green 1998), all of which express a single type of receptor protein (Buck and Axel 1991). There are ~ 1000 glomerular types, corresponding to the number of receptor cell types. The primary dendritic branches of a single mitral cell in the olfactory bulb lie largely within one glomerulus, and many different mitral cells share a single glomerulus (Shepherd and Green 1998). Here we will call the different mitral cells of a glomerulus the "glomerular repertoire" corresponding to a receptor type (Figure 3a). We propose that during odor recognition, the many neurons within a repertoire are all similarly driven by activation of their receptor type, but in addition have different positive bias currents driving them (Fig. 3b). We achieved this diversity here by assigning a random bias current to each repertoire cell, the magnitude of which was then assumed to be a fixed property of the cell.

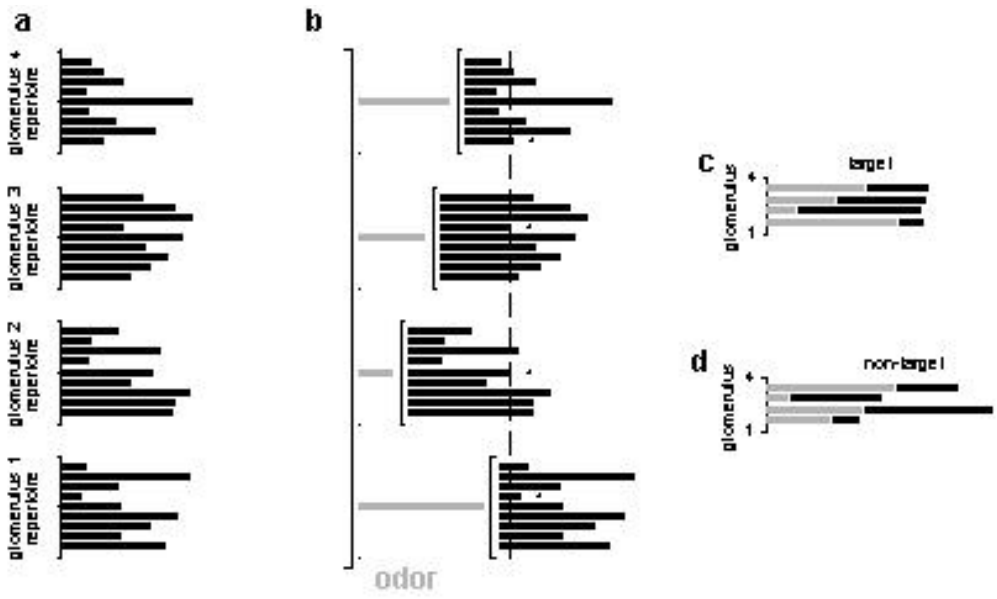

Figure 3. Glomerular repertoire and many-are-equals odor recognition. a, Each glomerulus has a set of repertoire cells, each of which is driven by a random bias current, with strength represented by the length of a black bar. b An odor that activates the 4 illustrated glomeruli above threshold. The length of the gray bars indicates the degree of activation of each receptor type. The *total* input current to each of the repertoire cells is its bias current plus its sensory current, and is represented by the location of the right ends of the black bars. We can find a set of repertoire cells, designated by asterisks, that all have total input current near to the same value (vertical dashed line); these can be used for recognition of this particular odor through the MAE operation. c. Sum of receptor activation plus selected repertoire cell bias currents (asterisked bars in b) for a target odor. d. Sum of receptor activation plus selected bias currents for a non-target odor.

Now, consider a particular odor that generates an analog pattern of glomerular activation. If each glomerulus has a repertoire of 'mitral' cells driven by random bias currents as in Fig 3 a, we can find a set of repertoire cells (asterisks in Fig. 3b) that receives, in addition to the receptor drive, a bias current such that the sum of the receptor and bias currents is similar, across glomeruli. Presentation of the target odor will then lead to similar net activation across the selected cells (Fig. 3c), while a different, non-target odor would lead to quite diverse activations across the selected cells (Fig. 3d). In essence, the set of bias currents of the selected mitral cells acts a "lock" that corresponds to the "key" of receptor activations for a specific (target) odor. The problem of odor recognition has thus been transformed into an MAE problem: the target odor is deemed present when many of the selected mitral cells have net activations (receptor + bias) that are closely similar to each other.

How does receptor activation depend on odor concentration? Let i index receptor types, and let the coverage r_i of receptor type i, in the presence of odor o, be given by

$$r_i = c_o k_{oi}$$

where c_o indicates the concentration of odor o and k_{oi} is a constant that depends on both receptor type and odor identity. The response defines a ~1000-dimension **odor vector having components r_i. The net signal reaching the mitral cells in the**

glomerulus corresponding to receptor type i will be denoted s_i. We assume that s_i encodes odor concentration in a roughly logarithmic fashion, that is:

$$s_i = k \log(1 + c_o k_{oi}/\theta)$$

$$s_i \approx 0 \quad \text{for} \quad c_o < \theta/k_{oi} \qquad s_i \approx k \log(k_{oi}/\theta) + k \log(c_o) \quad \text{for} \quad c_o > \theta/k_{oi} \qquad (1)$$

where k is the same for all receptor types. Below we will loosely refer to s_i as the glomerular activation corresponding to receptor type i. The term θ/k_{oi} represents a threshold concentration for a logarithmic transformation, below which glomerulus i is negligibly activated, and above which the activation of glomerulus i is approximately logarithmic. Thus a change in the concentration of an odor will lead to an additive change in s_i across all above-treshold glomeruli. Figure 4a-d shows the receptor + bias activations for a set of mitral cells driven by their target odor at different concentrations. At the highest concentrations (Figs. 4a,b), a change in concentration leads only to a change in the common level at which the mitral cells are driven: similarity of net drive across mitral cells is preserved on concentration changes. Thus, odor recognition based on the MAE operation will be concentration invariant. As concentration falls further, some receptors fall below threshold (Fig. 4c); corresponding mitral cells cease having a net drive similar to the others in the selected set. Since the MAE operation is robust to a few outliers, correct odor recognition will still be possible. At very low concentrations, where most or all receptors are below threshold, the mitral cells are driven only by their bias currents, and recognition is no longer possible (Fig. 4d).

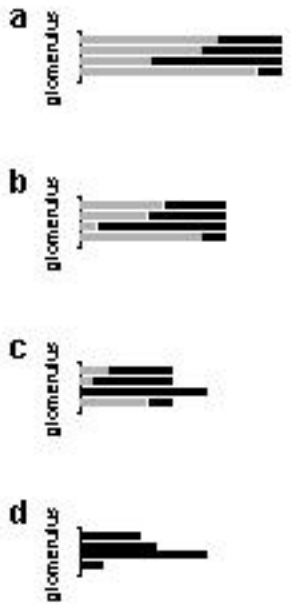

Figure 4. Activation of selected mitral cells for a target odor when that odor is presented at different concentrations. Same conventions as in Fig. 3a. Panels a, b, c, d, illustrate presentation of the target odor at successively lower concentrations. At the very lowest concentration (panel d), none of the

illustrated gomeruli are driven above threshold, and the mitral cell activities reflect only their bias currents.

Following Hopfield (1999), we have developed this specific description in the context of logarithmic encoding of concentration. However, the approach is generalizable to other situations. Let the activation of each receptor type r_i be given by some arbitrary but invertible function f_i that depends on both odor identity and odor concentration.

$$r_i = f_i(c_o ; o)$$

The logarithmic encoding described above is a special case of this, where odor identity determines threshold and k_{oi}, and odor concentration determines activation. To test the hypothesis that some specific odor o is present, the functions f_i can be inverted to get an estimate, from each r_i, of the concentration of odor o.

$$c_{oi} = f^{-1}(r_i ; o)$$

When odor o is present, the estimates c_{oi} will be the same across all receptors i. In contrast, when receptors are driven by a different odor, u, the estimates c_{oi} will be different to each other. This is true independently of concentration. Thus, an MAE operation on the estimates c_{oi} will result in concentration-invariant odor recognition.

The special case of logarithmic odor encoding permits concentration-invariance without inverting the function f(). There is a range over which this may correspond to the biological situation (Duchamp-Vieret et al. 2000; Meister and Bonhoeffer 2001). However, for some other odor encodings, neural implementation of the inversion of f() is also feasible. For example, if receptor activation depended linearly on odor concentration, $r_i = k_{i,odor}\, c_{odor}$, a variety of multiplicative synaptic weights from a receptor to the glomerular repertoire would replace the additive bias currents used in the logarithmic encoding case. To recognizing odor o, the mitral cell selected in glomerulus i would be the one that had synaptic weight most closely proportional to $1/k_{io}$. If odor encoding were quadratic, appropriate synaptic depression from receptors to repertoire cells could be used to invert f(); if odor encoding were proportional to the square root of concentration, synaptic facilitation could be used to invert f(). In the absence of detailed knowledge, we use here the simplest assumption, logarithmic encoding, but the essence of the computation can readily be implemented biologically for other encodings as well.

So far, we have described *concentration-invariant* **odor** *recognition***. Odor** *segmentation* **can be achieved by noticing that odors typically activate strongly only a subset of the available receptor types. Even if a strong distractor odor drives most of the target odor's receptor types above threshold, when both the distractor and the target are present, the relative activations of a significant fraction of the target's receptor**

types will be dominated by the target odor, allowing recognition of the target odor separately from the background. We explore this point more fully in the simulations of Fig. 6b below.

Simulation of a network of neurons for the model olfactory problem

We now demonstrate that the algorithm described above can be successfully implemented in a network of spiking neurons. Since we are unable to validate this implementation analytically, we instead performed extensive simulations using simple integrate-and-fire units. Our goal is to provide a proof-of-concept that noisy spiking neurons, using the MAE operation, can (1) achieve highly odor-selective synchronization; (2) do so in a manner invariant to a wide range of odor concentrations; and (3) do so in the presence of strong background distracting odors. In addition, we used the results of the simulations as the basis for evaluating what information about the odor vector is used in the network decision.

In our simulated network, there were 'odor inputs' from 400 glomeruli. Odors were modeled as in Hopfield (1999), with components of random strength spread over a relative intensity range of 10^6. An odor strength of 1.0 was defined as the strength of odor such that half (statistically) of the glomeruli received input which would drive them above threshold. For glomeruli driven above threshold, the input to the glomerulus was taken to be logarithmically encoded.

Each glomerulus had its own ensemble of repertoire cells, with a range of steady bias currents. The bias currents for repertoire cells were chosen at random within a defined range (Fig. 3a). For most of the simulations, 14 repertoire cells were assigned to each glomerulus.

All repertoire cells were also driven by a common oscillatory input current at 35 Hz. The total input to each repertoire cell was thus

$$I_{sensory} + I_{bias} + I_{periodic} \cos(\omega t)$$

where $I_{sensory}$ for each repertoire glomerulus was linearly proportional to s_i of equation (1). The repertoire cells were modeled as single compartment integrate-and-fire neurons with 20 ms. time constants. The cells fired 'action potentials' of negligible duration when their cell potential rose above a threshold. The cell potential was then reset to a lower level, where it was held for a 2 ms. dead time. Gaussian random noise was injected at each time step to each cell. Model cells with these same parameters were used in an earlier study (Hopfield and Brody 2000, 2001).

The odor stimulus, represented by the set of activations s_i, was a single 'sniff' lasting 0.5 sec. Mammals often sniff more rapidly than this, but also use multiple sniffs in making decisions. 0.5 sec. was chosen in order to simulate a sniffing phenomenon,

while getting enough information in a single sniff to make multiple sniffs unnecessary. When two odors are present, their components add. For the most general case dealt with in the simulations, described as being due to two odorants A and B simultaneously present, at concentrations c_A and c_B, the functional form of the coverage of each receptor type i is half a sine wave

$$r_i = [\sin(2\pi *(t-t_{start})] [c_A k_{Ai} + c_B k_{Bi}] \quad \text{if} \quad t_{start} < t < t_{start} + 0.5, \ 0 \text{ otherwise}$$

<----------------->
0.5-sec long sniff

where k_{Ai} is the input to the glomerulus from odor A alone at strength unity, and lies in the range 10^{-3} to 10^3 (and similarly for B). The activation s_i of glomerulus i is $\log(r_i)$ if $r_i > 1$ and is zero otherwise. Thus, if a single odor is present at concentration 1.0 ($c_A = 1$, $c_B=0$), half the glomeruli will be activated above threshold.

Synchrony across chosen subsets of repertoire cells (e.g., starred repertoire cells in Fig. 3c) is the event that indicates recognition of a specific odor. We use 'reporter' γ-cells, modeled (as earlier) in the same fashion as repertoire cells, but with faster (6 ms.) time constants, as simple detectors of the presence of synchrony. Connections to a γ-cell designed to recognize odor A were chosen by the following procedure: approximately 200 glomeruli are driven above threshold by odor A at concentration 1.0. One repertoire cell was chosen from each of these glomeruli to make a connection to the γ-cell; the repertoire cell was chosen so that the sum of the bias current and the sensory input from odor A at concentration 1.0 was at the center of the range of input currents where the repertoire cells showed good phase-locking. Synaptic currents in the γ-cells were generated as a result of the activity of the repertoire cells from which they received connections.

It is unlikely that biology uses single cell responses to correspond to single odors. We use a 'grandmother cell' representation of an odor as a surrogate of a 'highly selective cell' in neurobiology. In the γ-cell representation, there is no limit to the number of odors which could be separately recognized, each by its own γ-cell. In actuality, biology is much more likely to find a combinatorial representation of odors.

The γ-cells are driven by both excitatory and inhibitory synapses from the designated repertoire cells. The important part of this drive is excitatory input, which was modeled as a fast current, decaying exponentially with a 2 msec time constant. All excitatory synapses were of the same strength. Since the input to the repertoire cells does not involve the synapses in question, all synapse properties are irrelevant to the basic synchronization phenomenon.

Inhibitory synapses with a 6 msec time constant were introduced to help keep the system γ-cell inputs in automatic balance. The function served by the inhibition would be as well served by an inhibitory network driven by the inputs of all repertoire cells, and inhibiting all γ-cells. The strength of the inhibitory synapses

was chosen so that the total charge that flowed into a γ-cell due to an action potential in one of its repertoire cells was zero.

Results of Simulations

After constructing a network as described in the previous section, we recorded the behavior of the 'reporter' γ-cells in response to odors at various concentrations and odor mixtures. We now show that a γ-cell's spiking is highly odor-specific, responding to its target odor (only) over a wide range of concentrations, and recognizing its odor in the presence of strong distractor scents. We examine the extent to which the relative activations of different glomeruli, rather than merely the binary pattern of active vs. inactive glomeruli, are essential to these results and other olfactory tasks.

Figure 5b illustrates the behavior of a γ-cell designed to recognize an odor labeled A, in response to a presentation of A at concentration 1.5. The γ-cell spikes robustly. In contrast, in response to a different odor, B, at a concentration of 3.0, the γ-cell does not fire any spikes (Fig. 5d). Panel 5a illustrates the response of the γ-cell to a much higher concentration of odor A; panel 5c illustrates the response to a much lower concentration. Recognition events defined as 4 or more spikes from the A-selective γ-cell will thus be invariant to at least a 50-fold range of concentrations. We have not investigated the ultimate odor selectivity of the system here, but empirically we examined the response of an odor-selective γ-cell to 3,000,000 non-target random odors of strength 1.5. For computational efficiency, we used smaller, and therefore less selective, systems (the selectivity increases rapidly with the number of glomeruli). At 280 glomeruli, none of the $3*10^6$ random odors examined generated even a single γ-cell spike in response. A 400 glomerulus system would be even more selective. Requiring 4 spikes for odor identification is thus an extremely conservative criterion, given the unresponsiveness of the A-selective cell to random odors.

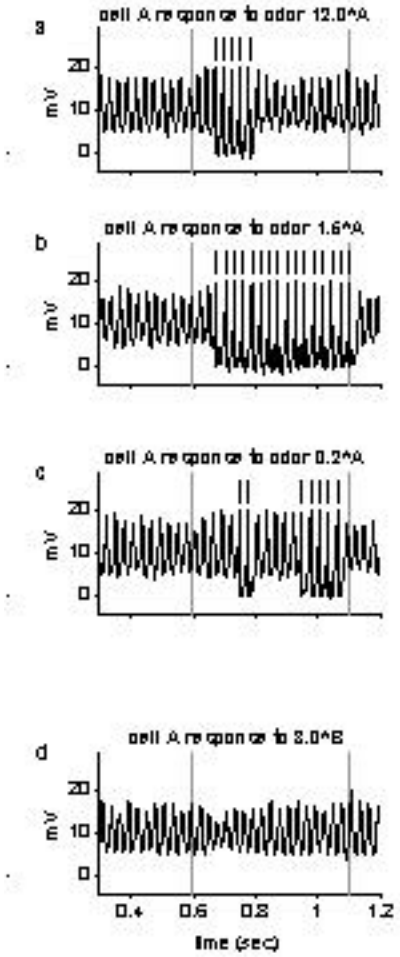

Figure 5. Odor-selective, concentration-invariant responses of a γ-cell. a-c, response of a γ-cell designed to be selective for odor A, upon presentation of odor A at three different concentrations. The solid line represents the membrane potential of the γ-cell. Spiking threshold is 20 mV; small vertical lines above the membrane potential indicate spikes. Vertical gray lines indicate the beginning and end of the 0.5-second long odor presentation. d, Response of same γ-cell upon presentation of a different, random, odor B.

Figure 6a shows the response of a γ-cell selective to odor A when presented with a mixture of odor A at concentration 1.0 and odor B at concentration 3.0. This cell, which does not respond to B alone, does respond briskly to the mixture. By contrast, a cell which is responsive to a third random odor C will not respond to A, B, or to this mixture (not shown). Similarly in Figure 6b the γ-cell selective for B responds to this mixture. These patterns indicate that a cell can identify its own target odor even in the presence of a background which is stronger than the target odor, and that this response is indeed the result of the presence of its target, not a general response of all cells to a complex mixture.

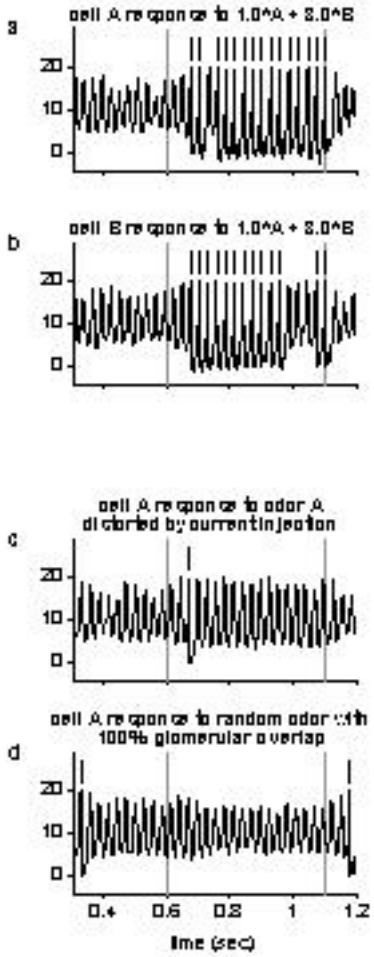

Figure 6. Odor selectivity and odor segmentation. Same format as Fig. 5. a, response of a γ-cell selective for odor A upon presentation of its target odor plus a stronger background odor. b, response of a γ-cell selective for the background odor B. c, response of γ-cell selective for odor A upon presentation of a distorted odor A (see text). d, response of γ-cell selective for odor A upon presentation of a random odor that has 100% overlap with odor A of glomeruli driven above threshold.

The analog strength of the glomeruli driven above threshold is essential in these recognitions. They are not merely the result of using the pattern of glomeruli driven above threshold (which varies strongly with concentration, Meister & Bonhoeffer 2001). To demonstrate this, we constructed an odor A' which drove above threshold precisely the same set of glomeruli driven above threshold by odor A, but with different, random analog activations. Figure 6d shows that the A-selective γ-cell did not fire any spikes in response to this `scrambled' odor A. The binary pattern of above-threshold glomeruli does not describe the selectivity of this system: the *analog* **pattern of activations is crucial.**

Unlike a response based only on the identities of glomeruli driven above threshold, *increasing* **the strength of the drive to above-threshold glomeruli can** *decrease* **the response of a γ-cell. Fig 6c shows the behavior of an odor A-specific γ-cell in response to odor A plus the injection of a common excitatory current into half of**

odor A's glomeruli. This current was equivalent to increasing the 'effective concentration' driving the injected glomeruli by a factor of 4. The result is that the γ-cell virtually ceases to generate action potentials, even though half of odor A's glomerular inputs were reinforced. This counterintuitive result comes about not by activation of an inhibitory pathway, but from the failure of all cells to share in a common phase of synchronization.

The analog aspect facilitates the analysis of mixtures (or rejecting backgrounds), an important olfactory task. Figure 6 indicates that two random odors A and B that activate only partially overlapping sets of glomeruli can be individually recognized in the mixture 1*A + 3 *B. We have also examined this separation problem in the much more difficult case, when odors *A and B are chosen so that they drive exactly the same glomerular set***. The decomposition is nonetheless successfully carried out by the network. The γ-cell for odor A responds somewhat less robustly than in the completely random case, but typically produces ~ 5 spikes in response to A alone and to 1*A + 3*B, yet produces zero spikes in response to B alone (data not shown). In this case, all three of A alone, B alone, and 1*A + 3 *B have exactly the same pattern of glomeruli driven above threshold. Odor recognition based only on the pattern of glomeruli activated above threshold cannot accomplish this task, nor can it explain the capacity of the network to deal with this problem.**

The quantity of analog information used

For channels that are driven strongly enough that their responses are above threshold, the pattern of *analog* **responses is essential to the recognition process. How much information is carried by the analog degree of activation of each glomerulus? We address this quantitatively by varying the number of repertoire cells per glomerulus: if we had a system with only 1 repertoire cell per glomerulus, then for glomeruli driven above threshold by an odor that single repertoire cell would be the 'nearest' one in the selection process used to choose connections for recognizing the odor. In this case, therefore, the representation of known odors would have no information about the relative degrees to which different glomeruli were driven above threshold by the target odor—in this limit, there would be zero analog information stored. When there are n repertoire cells per glomerulus, knowing which is nearest contains some information about the analog strengths of the target odor; the larger the number of repertoire cells n, the more precisely those analog strengths are known.**

The information in 'bits' due to knowing which repertoire cell (of the random set) is closest can be evaluated in straightforward but tedious mathematics. We have done a Monte-Carlo evaluation of that mathematics to show that the information available in principle per glomerulus for n=2, 4, 7, 14, and 28 repertoire cells is 0.87, 1.76, 2.53, 3.50 and 4.48 bits, respectively. For large n this grows as 1 bit per two-fold increase in n. How much of this information is used by the system?

Fig 7 shows histograms of the number of spikes produced by a γ-cell nominally selective to an odor when that odor is present, as a function of n, the number of repertoire cells per glomerulus available to the system. With 2 repertoire cells, the γ-cell does not respond at all: the representation of the analog strengths of the target odor is too poor to represent the odor within the accuracy detectable by the system. As more repertoire cells are added, the analog strengths are better and better represented by the choice of which repertoire cells are connected to the γ-cell, and more spikes are likely to be generated. The response saturates shortly beyond 14 repertoire cells. This shows that the system does not use a storage precision for known odors greater than the information per glomerulus corresponding to 14 repertoire cells/glomerulus. Thus, f
rom Fig. 7 we conclude that the useful precision in the analog value of the logarithm of the odor intensity is about 3.5 bits.

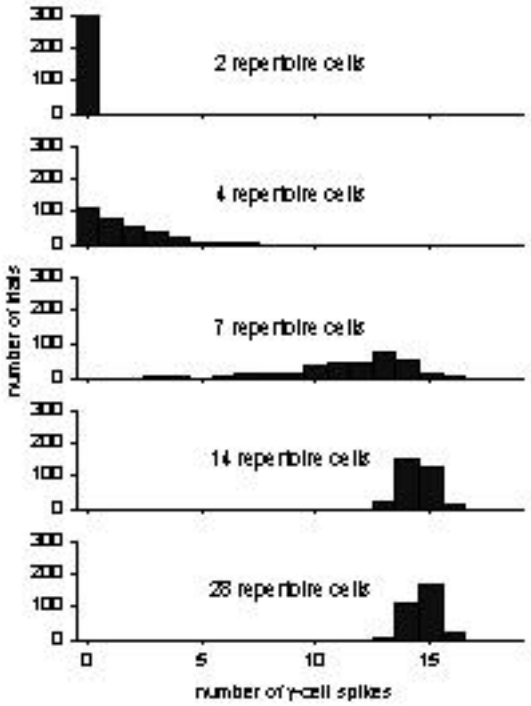

Figure 7. The dependence of a γ-cell designed to be responsive to a given odor on the number of repertoire cells used. Statistics collected on 300 odors, with 2, 4,7, 14, and 28 repertoire cells per glomerulus, representing 0.87, 1.76, 2.53, 3.50, and 4.48 bits of information about the degree to which a glomerulus is driven above threshold. N_glomeruli = 400.

To further pursue the selectivity available through making use of analog values, we examined the behavior of a γ-cell selective for an odor A in response to a set of random odors *iall involving only exactly the same glomeruli as odor A being driven above threshold***. Thus, the only thing that distinguishes these odors from each other is their different analog strengths across the set of activated glomeruli. If by 'odor pattern' one were to mean only the identity of the set of glomeruli that are driven above threshold, all these random odors would be functionally equivalent to odor A, and**

they should all produce the vigorous ~ 15 spike response that odor A does. We applied 3,000,000 such random odors to a system which, for computational efficiency, was a smaller system with 200 glomeruli. Only a small fraction, $4*10^{-6}$ of the total produced a single spike, and *none produced more than one spike.* **The selectivity and signal to noise increase with the number of glomeruli; with 280 glomeruli, no action potentials were generated by any of the $3*10^{6}$ random odors. With 400 or more glomeruli, the** *functional* **selectivity due to analog information will be enormous. In the more natural case of random odors across all glomeruli (i.e., not always using the same set of glomeruli activated above threshold), the functional selectivity is even greater.***functional*

Conclusion

We have shown that noisy spiking neurons, using the many-are-equal operation, can be configured, in a manner reminiscent of the anatomical and functional organization of the olfactory bulb, to solve some of the major computational problems faced by olfactory systems in a natural environment. The many mitral cells for each glomerulus are tuned to the same odors, but they provide a repertoire of different response strengths that are the basis for representing, and computing with, analog signal strength. The presence of target odors is indicated by synchronization across a number of 'mitral' cells in different glomeruli, and is highly odor-selective. This selective synchronization is invariant to a 50-fold change in concentration. Odor-selective synchronization still occurs in the presence of strong background odors, thus allowing odor segmentation and the decomposition of mixtures of known odors into their components.

The MAE operation allows a system to ignore badly contaminated information, as long as it does not affect too large a fraction of information channels. The use of MANY, rather than ALL, is essential to carrying out olfactory tasks, for a particular odorant can easily dominate many glomeruli while other glomeruli are responding to other odorants which are simultaneously present. In the simulations, the higher the number of similarly-activated mitral cells connected to a γ-cell, the more strongly they will drive the γ-cell. We can choose connections strengths and thresholds for the γ-cell such that strong firing will result if half or more of the mitral cells are synchronized. This allows high odor selectivity, since it is highly unlikely that a random odor would match a subset of this size. At the same time, this threshold level allows significant robustness, since any half of the receptor set may be corrupted without disrupting odor recognition (Hopfield *et al***. 1998).**

In this scheme, it is the relative activations between glomeruli that are the key to the definition of a target odor (Fig. 3a). Thus, differentiating between two individually presented random odors that have as much as 100% overlap in the identity of activated receptors is readily done (compare Figs. 4a-c to Figs. 5a-b; in the situation more nearly resembling the natural olfactory problem, two odors have much less glomerular overlap, further simplifying the problem). More surprisingly, and thanks to the robustness to corrupted components of the MAE operation, the system can even separately recognize two simultaneously presented random odors with

100% glomerular overlap. Such problems would be totally intractable if odors were described merely by a 'spatial pattern,' i.e., by the identity of the glomeruli excited above some fixed threshold.

References

Buck, L. and Axel, R. (1991). A novel multigene family may encode odorant receptors--a molecular-basis for odor recognition. Cell *65*, 175-187.

Buck L. B. (1996). Information coding in the vertebrate olfactory system. Ann. Rev. Neurosci *19*, 517-544.

Duchamp-Vieret, P., Duchamp, A., and Chaput, M. A. (2000). Peripheral Odor Coding in the Rat and Frog: Quality and Intensity Specification, J. Neurosci. *20*, 2383-2390.

Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. Nature *376*, 33-36.

Hopfield J. J. (1996). Odor space and olfactory processing: Collective algorithms and neural implementation. Proc. Nat. Acad. Sci. *96*, 12506-12511.

Hopfield, J. J., Brody, C. D., and Roweis, S. (1998). Computing with action potentials. Adv. in Neural Inf. Processing *10*, 166-172.

Hopfield, J. J. and Brody, C. D. (2000). What is a moment? "Cortical" sensory integration over a brief interval. Proc. Nat. Acad. Sci. (USA) *97*, 13919-13924.

Hopfield, J. J. and Brody, C. D. (2001). What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. Proc. Nat. Acad. Sci. (USA) *98*, 1282-1287.

Johnson, D. H. (1980). The relationship between spike rate and synchrony in responses of auditory-nerve fibers to single tones. J. Acoustical Soc. Am. *68,* 1115-1122.

Lavine, R. A. (1971). Phase locking in response of single neurons in cochlear nucleus complex of the cat to low frequency tonal stimuli. J. Neurophysiol. *34*, 467-483.

Meister, M. and Bonhoeffer, T. (2001). Tuning and topography in an odor map on the rat olfactory bulb. J. Neurosci. *21* 1351-1360.

Pouille, F. and Scanziani, M. (2001) Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. Science *293*, 1159-1163.

Shepherd, G. M. and Green, C. A. (1998). Olfactory Bulb, in The Synaptic Organization of the Brain, Shepherd, G. M., ed., Oxford, New York, 159-204.

Sicard, G. and Holley A. (1984). Receptor cell responses to odorants--similarities and differences among odorants. Brain Res. *292*, **283-296.**

Swadlow, H. A. (2001). The imact of 'bursting' thalamic impulses at a neocortical synapse. Nat. Neurosci. *4*, **402-408.**